

AMENDMENTS TO SPECIFICATION:

Please replace paragraph [017] with the following amended paragraph:

[017] Figure 1 is a diagram illustrating an example network communication system 100 configured in accordance with one embodiment of the systems and methods described herein. Network communication system 100 comprises a source device 101 interfaced with a plurality of destination devices ~~111a, 111b, 111c, 111d and 111e~~ via connection interfaces 103a, 104a, and 105a and communication connections 106-110. Source device 101 comprises a source data buffer 102 that comprises data to be replicated on each destination device ~~111a-111e~~, e.g., in an associated destination data buffer ~~117a, 117b, 117c, 117d and 117e~~. Source device 101 can also include one or more synchronization mechanisms, such as synchronization mechanisms 103, 104, and 105, as well as a cluster manager 118. The source data buffer 102 can be coupled directly or indirectly to synchronization mechanisms 103, 104 and 105.

Please replace paragraph [018] with the following amended paragraph:

[018] The number of destination devices ~~111a-111e~~ and the number of associated communication connections are shown for illustration only. Only a small number of connections and destination devices ~~111a-111e~~ are shown for clarity and should in no way imply a limit or suggestion as to the number of communication devices ~~111a-111e~~ that can be supported using the systems and methods described herein.

Please replace paragraph [019] with the following amended paragraph:

[019] Often, communication connections 106-110 can have different performance capabilities. As described below, performance capability can be defined in a variety of ways, such as the bandwidth capability of each connection. Thus for example, connections 106 and 107 can have a relatively high bandwidth capability, while connection 108 can have a somewhat lower bandwidth capability and connections 109 and 110 can have a still lower bandwidth capability. Accordingly, cluster manager 118 can be configured to group destination devices ~~111-111a-111e~~ or communication links 106-110, into performance clusters, e.g., performance clusters 119, 120, and 121, based on their similar performance capabilities. In Figure 1, performance cluster 119 is shown as a high rate or high performance cluster; performance cluster 120 as shown as a medium rate or intermediate performance cluster; and performance cluster 121 is shown as a low rate or low performance cluster. The number of clusters needed may vary as system requirements dictate. The cluster manager can be further configured to assign a synchronization mechanism 103, 104, or 105 to each of the performance clusters 119, 120, and 121. Synchronization mechanisms 103, 104, or 105 can be configured to then send updates to associated destination devices ~~111-111a-111e~~ or associated communication links 106-110 in a manner that is optimized for the performance capabilities of the associated destination devices or communication links 106-110.

Please replace paragraph [020] with the following amended paragraph:

[020] Source device 101 can be configured to share data stored in the source data buffer 102 with the plurality of destination devices ~~111-111a-111e~~. Communication connections 106-

110 provide mechanisms for transferring data, i.e., physical communications channels, while the synchronization mechanisms 103-105 and corresponding synchronization mechanisms ~~116~~ 116a, 116b, 116c, 116d and 116e associated with destination devices ~~111~~ 111a-111e can be configured to provide the computation and protocols needed to share data between source data buffer 102 and destination data buffers ~~117~~ 117a-117e over communication connections 106-110.

Please replace paragraph [022] with the following amended paragraph:

[022] In one embodiment, the similarity in performance capability of the various communication connections is determined, in step 304, by maintaining statistics for data rate capability of each communication connection. Such statistics can be derived, for example, from observing data transfers between source device 101 and destination devices ~~111~~ 111a-111e. In another embodiment, connection security associated with each of the communication connections 106-110 can be used to determine similar performance capabilities in step 304. In still another embodiment, the error rate associated with data transfer of each communication connection can be used in step 304. In yet another embodiment, latency associated with data transfer of each communication connection can be used in step 304. In fact, it will be understood that a variety of performance parameters and related information can be used to determine the similarity in the performance capabilities of the various communication connections. Thus, nothing within the specification or the claims that follow should be seen as limiting the systems and methods described herein to the use of any particular parameters or set of parameters.

Please replace paragraph [023] with the following amended paragraph:

[023] In one embodiment, the number of synchronization mechanisms can be determined dynamically and can change as needed to accommodate destination devices ~~111~~ 111a-111e as they connect with source device 101. Several algorithms can be used in selecting the number of synchronization mechanisms, some of which are described below. This dynamic capability can allow for a trade off between improved client service, which results when there are fewer destination devices ~~111~~ 111a-111e per cluster, and reduced server resource usage, which results from having a large amount of clusters. Thus, for example, if there are only a few destination devices ~~111~~ 111a-111e, or if client service is important, then cluster manager can assign, for example, each destination device ~~111~~ 111a-111e to its own synchronization mechanism. On the other hand, if there are a lot of destination devices ~~111~~ 111a-111e, or if client service is not as important, then cluster manager can assign fewer synchronization mechanisms. Moreover, depending on the embodiment, cluster manager 118 can be configured to dynamically update the destination device groupings and add or remove synchronization mechanisms as required.

Please replace paragraph [024] with the following amended paragraph:

[024] Further, in one particular embodiment, the correspondence between a destination device ~~111~~ 111a-111e and a particular synchronization mechanism 103, 104, or 105 can also be dynamic. In other words, the corresponding communication connection for a particular destination device ~~111~~ 111a-111e can be moved to a different synchronization mechanism if the corresponding performance capabilities change such that a different performance cluster 119,

120, or 121 is more appropriate. Thus, for example, cluster manager 118 can be configured to monitor, in step 312, a set of statistics associated with the performance of each communication connection 106-110 and to detect any change therein. If a significant change is detected, then the statistics can be used to determine if another performance cluster 119, 120, or 121 is more appropriate for the particular destination device ~~111 111a-111e~~.

Please replace paragraph [025] with the following amended paragraph:

[025] In one embodiment, all connections 106-110 or all destination devices ~~111 111a-111e~~ begin a session as part of a primary performance cluster. After a small number of updates the average latency for each destination device ~~111 111a-111e~~ or communication connection 106-110 is gathered. Cluster manager 118 can be configured to then perform a cluster division (step 306) to organize destination devices ~~111 111a-111e~~ according to their performance levels, e.g., their average latencies. For example, an initial calculation of the average latencies for each of the plurality of connections can be performed and used to determine the mean latency for the primary performance cluster. A standard deviation relative to the mean can also be calculated. The number of performance clusters required can then be determined based on the percentage of communication connections 106-110 with latencies within a certain number of standard deviations from the mean.

Please replace paragraph [027] with the following amended paragraph:

[027] Further, in one embodiment, communication connections 106-110 can be placed into an appropriate performance cluster (step 308) using an algorithm, such as the K-means

algorithm. The K-means algorithm is a partitioning method based on minimizing the sum-of-squares distance between average latencies for each communication connection 106-110 and the mean latency for the primary performance cluster, allowing division of communication connections 106-110 into (K) performance clusters. This is an iterative approach that maximizes the ratio of variation between performance clusters relative to the variation within a performance cluster. This approach allows a quick calculation with a resulting distinct separation of performance levels. Depending on the embodiment, the K-means algorithm is executed periodically or as needed, e.g., if there is a change in performance capabilities (step 312) or a new destination device ~~111 111a-111e~~ joins the session (step 314).

Please replace paragraph [029] with the following amended paragraph:

[029] A new insertion into a performance cluster can require a resynchronization for destination devices ~~111 111a-111e~~ within the performance cluster on the next cluster-wide update. In one embodiment, only the newly added destination device ~~111 111a-111e~~ need be resynchronized while the other destination devices ~~111 111a-111e~~ remain synchronized.

Please replace paragraph [030] with the following amended paragraph:

[030] Figure 3 is a flow chart illustrating an exemplary method for updating destination data buffers ~~117 117a-117e~~ associated with a particular performance cluster 119-121 in accordance with one embodiment of the systems and methods described herein. Thus, in step 402, an update can be sent using a synchronization mechanism 103, 104, or 105. When the update is sent in step 402, a timer can then be started, in step 404. In step 406, the associated

communication connections are monitored until one of the associated destination devices, requests another update. When another request is received the timer is stopped in step 408. The timer value can then be used to determine the latency associated with the communication connection for the requesting destination device ~~111 111a-111e~~. For each destination device ~~111 111a-111e~~ two metrics can be determined and maintained in step 410. The two metrics can include a “Total Session Average Latency” and a “Recent Average Latency.” The “Total Session Latency” can be an average of all latency values associated with the requesting destination device ~~111 111a-111e~~. The “Recent Average Latency” can be an average of some number of the most recent latency values for the requesting destination device ~~111 111a-111e~~.

Please replace paragraph [031] with the following amended paragraph:

[031] A destination device ~~111 111a-111e~~ that does not respond within a timeout threshold, as determined in step 412, can be removed from its performance cluster, in step 414, so that other destination devices ~~111 111a-111e~~ in the performance cluster can still receive updates. This can, for example, prevent a network interruption or an issue associated with a destination device ~~111 111a-111e~~ from harming the other cluster participants’ experiences. If a destination device 111 has reached this timeout, in step 412, but eventually responds, in step 416, then it can still be allowed to receive full-buffer updates, in step 418, e.g., until its Recent Average Latency performance merits insertion back into one of the performance clusters 119-121.

Please replace paragraph [032] with the following amended paragraph:

[032] In one embodiment, the synchronization mechanisms 103-105 and the synchronization mechanisms ~~116~~ 116a-116e can be configured to operate by dividing the data in source data buffer 102 into a number of blocks or sections. Initially, or whenever synchronization is lost, a complete set of all blocks can be sent from source device 101 to a destination device ~~111~~ 111a-111e. The associated destination device buffer ~~117~~ 117a-117e can then be updated using the complete set of blocks so that it is a replica of the source data buffer 102. Subsequently source device 101 can send only blocks that have changed subsequent to the last update sent to the destination data buffer ~~117~~ 117a-117e. This approach can, for example, result in considerable savings in network bandwidth.

Please replace paragraph [034] with the following amended paragraph:

[034] It should be noted that the original source data buffer can be located on a remote source device, e.g., a device that is not immediately connected with destination devices ~~111~~ 111a-111e. For example, figure 4 is a diagram illustrating an example network communication system 200, which comprises a remote source device 205, with remote source data buffer 207, in accordance with one embodiment of the systems and methods described herein. Remote source device 205 can be interfaced with destination devices ~~111~~ 111a-111e via an intermediate source device 201. Thus, intermediate source device 201 can comprise a cluster manager 118, which can be configured to group communication connections 106-110 into performance clusters 119-121 using synchronization mechanisms 103-105.

Please replace paragraph [035] with the following amended paragraph:

[035] In addition, intermediate source device 201 can comprise an intermediate source data buffer 202, which can be kept in synchronization with remote source data buffer 207 using synchronization mechanisms 203 and 205. Destination data buffers ~~117~~ 117a-117e can then be kept in synchronization with intermediate source data buffer 202 as described above in relation to source data buffer 102.

Please replace paragraph [036] with the following amended paragraph:

[036] In one example of figure 4, synchronization mechanism 103 can be configured to provide updates to performance cluster 119, which comprises the highest data rate destination devices ~~111~~ 111a-111e and communication connections 106 and 107. The update interval required by synchronization mechanism 103 can thus determine the highest update rate needed and can therefore also serve as the update interval used by synchronization mechanisms 203 and 205 for updating intermediate source data buffer 202. This approach can be used to avoid redundant or excessive data requests for communication connection 204.